arXiv:1612.00939v2 [stat.ME] 14 Feb 2017

Projection Sparse Principal Component Analysis: an efficient least squares method

Giovanni Maria Merola
Department of mathematical Sciences
Xi'an Jioatong Liverpool University
giovanni.merola@xjtlu.edu.cn

February 15, 2017

Abstract

We propose a new sparse principal component analysis (SPCA) method in which the solutions are obtained by projecting the full cardinality principal components onto subsets of variables. The resulting components are guaranteed to explain a given proportion of variance. The computation of these solutions is very efficient. The proposed method compares well with the optimal least squares sparse components. We show that other SPCA methods fail to identify the best sparse approximations of the principal components and explain less variance than our solutions. We illustrate and compare our method with the analysis of a real dataset containing socioeconomic data and the computational results for nine datasets of increasing dimension with up to 16,000 variables.

1 Introduction

Data dimensionality reduction methods approximate a set of variables on a lower dimensional space. Principal components analysis (PCA) is the oldest and most popular of these methods. Given the measurements of p variables (or features) on n subjects (or objects) the principal components (PCs) are the linear combinations of the variables that best approximate the data on an orthogonal subspace with respect to the least squares criterion (LS) (Pearson, 1901).

The PCs are difficult to interpret and not informative on important features of the dataset because they are combinations of all the observed variables, as already pointed out by Jeffers (1967). Sparse principal components analysis (SPCA) is a variant of PCA in which the components computed are combinations of just a few of the observed variables. The aim of SPCA is to produce a reduction of the dimensionality that is easy to interpret.

Formally, let $\mathbf{X}$ be an $(n \times p)$ matrix containing n observations on p mean centred variables. The information contained in the dataset is summarised by its *total variance* defined by $||\mathbf{X}||^2 = trace(\mathbf{S})$, where $\mathbf{S} = \mathbf{X}^\intercal\mathbf{X}$ is the covariance matrix (omitting a scaling constant which is irrelevant for our analysis).

A *component* is any linear combination of the columns of $\mathbf{X}$, defined by $\mathbf{t} = \mathbf{X}\mathbf{a}$, where the elements of the vector $\mathbf{a}$ are called *loadings*. The least squares estimates of $c < p$ components $\mathbf{T} = (\mathbf{t}_1, \ldots, \mathbf{t}_c)$ are obtained by minimising the squared (Frobenius) norm of the difference of the data matrix from its projection onto the components, $\widehat{\mathbf{X}}(\mathbf{T})$. That is, they are the solutions to

$$\mathbf{T} = \arg\min_{\mathbf{T}} ||\mathbf{X} - \widehat{\mathbf{X}}(\mathbf{T})||^2 = \arg\max_{\mathbf{T}} ||\widehat{\mathbf{X}}(\mathbf{T})||^2. \quad (1)$$

Further constraints are necessary in order to identify the individual components. The *variance explained*, $||\widehat{\mathbf{X}}(\mathbf{T})||^2$, maximised by the least squares criterion is used to measure the performance of the components in approximating the data.

The PCs are the mutually orthogonal components that sequentially explain the largest proportion of variance of the $\mathbf{X}$ matrix. The (vectors of) loadings of the PCs are the eigenvectors of the covariance matrix $\mathbf{S}$ and the variance explained by each PC is equal to its norm. For this reason the PCs are commonly defined as the components with unit norm loadings that sequentially have largest norm (Hotelling, 1933).

A sparse component is a linear combination of a subset of columns of the $\mathbf{X}$ matrix, or *block* of variables, $\dot{\mathbf{X}}$, defined by the loadings $\dot{\mathbf{a}}$ as $\mathbf{t} = \dot{\mathbf{X}}\dot{\mathbf{a}}$. The number of variables in the block is referred to as the *cardinality* of the loadings. Sparse components are not eigenvectors of the covariance matrix and the variance that they explain is not equal to their norm. Since, in most cases they are not required to be orthogonal, new, more general measures of their performance are needed.

In recent years a great many SPCA methods have been proposed (among others, Jolliffe and Uddin, 2000; Moghaddam et al., 2006; Zou et al., 2006; Sriperumbudur et al., 2009; Shen and Huang, 2008; Wang and Wu, 2012). These use different optimisation approaches to find sparse components that

have largest norm. Because of the adoption of this objective function the resulting components do not maximise the variance explained. Furthermore these methods cannot reproduce the sparsest expression of the PCs when the data is rank deficient and the components can be combinations of perfectly correlated variables, as we will show later.

In this paper we adopt the least squares sparse components analysis (LS SPCA, Merola, 2015) approach to derive sparse components. In this approach the components are derived by minimising the least squares estimation criterion under sparsity constraints. LS SPCA components explain the most possible variance for a given block of variables and do not suffer from the same drawbacks mentioned above for other methods.

Identifying the blocks of variables from which to compute the components is an intractable NP-hard problem (Moghaddam et al., 2006). Therefore, greedy approaches are used to solve the SPCA problem. The techniques used to find components with the largest norm used in conventional SPCA are not appropriate for the maximisation of the variance explained. Merola (2015) proposed a backward elimination procedure which gives excellent results but does not scale well to larger datasets.

In this paper we propose computing sequentially the sparse components by projecting each PC on blocks of variables that explain a predetermined proportion of their variance. We show that the total variance explained by these projections is not less than the variance of the PC explained by the variables in the blocks. This means that we can select the variables in the blocks with an economical supervised forward selection and then compute the sparse component with a simple orthogonal projection. We will refer to this method as projection SPCA.

We also show that the lower bound for the variance explained by the projection SPCA components is also valid for the LS SPCA components computed from the same blocks of variables. Furthermore, we provide a novel interpretation of the LS SPCA components showing that these are the PCs of the projection of the data matrix onto the current block of variables.

We propose a simple but efficient algorithm for computing projection sparse components. We will show that sparse components for datasets with over 16,000 variables can be computed in a matter of seconds with this algorithm. Empirically we find that the time complexity of the algorithm is slightly larger than $O(p^2)$.

The paper is organised as follows: in the next section we derive projection SPCA comparing it to LS SPCA and to conventional SPCA methods. In Section 3 we give computational details. In Section 4 we illustrate the usefulness of projection SPCA with an application on a real set of social data and show its performance on different large datasets also comparing it

with that of conventional SPCA. In Section 5 we give some final remarks.

2 Projection SPCA

In this section, for ease of exposition, we will refer to quantities computed in a generic step of an algorithm without the index identifying the order of the step; the index will be used only when needed.

We denote a generic vector of loadings as $\mathbf{a}$. When sparsity is included in the model the symbol $\dot{\mathbf{a}}$ denotes the vector of $c < p$ non-zero loadings corresponding to the variables in the block $\dot{\mathbf{X}}$. In some cases we write the sparse vector $\dot{\mathbf{a}}$ in its full cardinality expression as $\mathbf{a}$ where missing values are set equal to zero. Principal components will be denoted as $\mathbf{u}$ and their loadings as $\mathbf{v}$.

2.1 Relationship between variance explained and norm of the components

When estimating a set of components, the minimisation of the least squares criterion is equivalent to the maximisation of the variance explained, as shown in Equation (1). For a set of components $\mathbf{T} = (\mathbf{t}_1, \ldots, \mathbf{t}_c) = \mathbf{X}(\mathbf{a}_1, \ldots \mathbf{a}_c)$ this is the total variance of the projection of $\mathbf{X}$ onto $\mathbf{T}$. If we denote the projector onto $\mathbf{T}$ by $\mathbf{P_T} = \mathbf{T}(\mathbf{T}^\intercal\mathbf{T})^{-1}\mathbf{T}^\intercal$ so that this projection is $\widehat{\mathbf{X}}(\mathbf{T}) = \mathbf{T}(\mathbf{T}^\intercal\mathbf{T})^{-1}\mathbf{T}^\intercal\mathbf{X}$, the variance explained by the components in $\mathbf{T}$ is equal to

$$vexp(\mathbf{T}) = ||\widehat{\mathbf{X}}(\mathbf{T})||^2 = trace(\mathbf{X}^\intercal\mathbf{T}(\mathbf{T}^\intercal\mathbf{T})^{-1}\mathbf{T}^\intercal\mathbf{X}).$$

In PCA the variance explained is maximised simultaneously for all components, $\mathbf{U} = (\mathbf{u}_1, \ldots, \mathbf{u}_c) = \mathbf{X}(\mathbf{v}_1, \ldots, \mathbf{v}_c)$, under the constraint that these are mutually orthogonal. Because of these constraints, the variance explained by the c components can be broken down into the sum of the variances of the complement of $\mathbf{X}$ orthogonal to the previously computed components explained by each component. This complement is the difference of the data matrix from its projection onto the previous components $\mathbf{X} - \widehat{\mathbf{X}}(\mathbf{U})$ and is equal to $\mathbf{Q} = \mathbf{X} - \mathbf{U}(\mathbf{U}^\intercal\mathbf{U})^{-1}\mathbf{U}^\intercal\mathbf{X}$. Therefore, the variance explained by the next principal component $\mathbf{u} = \mathbf{X}\mathbf{v}$ is equal to

$$vexp(\mathbf{u}) = \frac{\mathbf{u}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\mathbf{u}}{\mathbf{u}^\intercal\mathbf{u}} = \frac{\mathbf{v}^\intercal\mathbf{Q}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\mathbf{Q}\mathbf{v}}{\mathbf{v}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\mathbf{v}},$$

where the second equality comes from the orthogonality constraints for which $\mathbf{u}^\intercal\mathbf{X} = \mathbf{u}^\intercal\mathbf{Q} = \mathbf{v}^\intercal\mathbf{Q}^\intercal\mathbf{Q}$. Setting the derivatives of $vexp(\mathbf{u})$ with respect to $\mathbf{v}$

equal to zero gives

$$\frac{dvexp(\mathbf{u})}{d\mathbf{v}} = \mathbf{Q}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\mathbf{Q}\mathbf{v} - \mathbf{Q}^\intercal\mathbf{Q}\mathbf{v}\lambda \equiv \mathbf{Q}^\intercal\mathbf{Q}\mathbf{v} - \mathbf{v}\lambda = \mathbf{0},$$

with $\lambda = vexp(\mathbf{u}) = \mathbf{v}^\intercal\mathbf{Q}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\mathbf{Q}\mathbf{v}/\mathbf{v}^\intercal\mathbf{Q}^\intercal\mathbf{Q}\mathbf{v}$. For the j-th PC the solution that maximises the variance explained is the first eigenvector of $\mathbf{Q}_j{}^\intercal\mathbf{Q}_j$, which is equal to the j-th eigenvector of $\mathbf{S}$. Therefore, $vexp(\mathbf{u}_j)$ simplifies to

$$vexp(\mathbf{u}_j) = \frac{\mathbf{v}_j{}^\intercal\mathbf{Q}_j{}^\intercal\mathbf{Q}_j\mathbf{Q}_j{}^\intercal\mathbf{Q}_j\mathbf{v}_j}{\mathbf{v}_j{}^\intercal\mathbf{Q}_j{}^\intercal\mathbf{Q}_j\mathbf{v}_j} = \frac{\mathbf{v}_j{}^\intercal\mathbf{Q}_j{}^\intercal\mathbf{Q}_j\mathbf{v}_j\lambda_j}{\mathbf{v}_j{}^\intercal\mathbf{v}_j\lambda_j} = \frac{\mathbf{v}_j{}^\intercal\mathbf{S}\mathbf{v}_j}{\mathbf{v}_j{}^\intercal\mathbf{v}_j} = \lambda_j, \qquad (2)$$

because $\mathbf{Q}_j{}^\intercal\mathbf{Q}_j\mathbf{v}_j = \mathbf{S}\mathbf{v}_j = \mathbf{v}_j\lambda_j$, where λ_j is the j-th eigenvalue of $\mathbf{S}$. This equation is the connection between Pearson's least squares definition of PCA and Hotelling's maximal norm definition.

Since the PCs loadings are eigenvectors of $\mathbf{S}$, the orthogonality constraints $\mathbf{v}_j{}^\intercal\mathbf{S}\mathbf{v}_k = 0, j \neq k$ are equivalent to the constraints $\mathbf{v}_j{}^\intercal\mathbf{v}_k = 0, j \neq k$. By adding the requirement that $\mathbf{v}_j{}^\intercal\mathbf{v}_j = 1$ Equation (2) simplifies to Hotelling's definition of PCA:

$$\begin{aligned} \mathbf{v}_j &= \arg\max_{\mathbf{v}_j} \mathbf{v}_j{}^\intercal\mathbf{S}\mathbf{v}_j \qquad (3) \\ &\quad \mathbf{v}_j{}^\intercal\mathbf{v}_k = \delta_{jk}, \end{aligned}$$

where $\mathbf{v}_j{}^\intercal\mathbf{S}\mathbf{v}_j = \lambda_j$ is the norm of the components with unit norm loadings (hereafter referred to simply as norm of the components) and δ_{ij} is the Kronecker delta..

The simplified expression of the variance explained by a component in Equation (2) is valid if and only if the loadings are eigenvectors of $\mathbf{S}$. Hence it is not valid for a sparse component $\mathbf{t} = \dot{\mathbf{X}}\dot{\mathbf{a}}$ because in this case the first order conditions are

$$\frac{dvexp(\mathbf{t})}{d\dot{\mathbf{a}}} = \dot{\mathbf{Q}}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\dot{\mathbf{Q}}\dot{\mathbf{a}} - \dot{\mathbf{Q}}^\intercal\dot{\mathbf{Q}}\dot{\mathbf{a}}\gamma = \mathbf{0},$$

which cannot be simplified as above and the solutions are not eigenvectors of $\mathbf{S}$. This should be expected because, in general, eigenvectors are not sparse. For this reason, the properties of the PCs are not valid for sparse components; in particular, the norm of the sparse components is not equal to the variance that they explain and the loadings are not required to be orthonormal (see Merola, 2015, for a more details).

2.2 Extra variance explained

Since sparse loadings are not eigenvectors of the covariance matrix and sparse components are not necessarily orthogonal, we need to introduce new measures for the variance explained.

Let $\mathbf{T} = (\mathbf{t}_1, \ldots, \mathbf{t}_c)$ be the matrix containing the first $c < p$ components computed. If a new component, $\mathbf{t}$, that enters the model is correlated with the previously computed components — that is, if $\mathbf{t}_i{}^\intercal \mathbf{t}_j \neq 0,\ i < j$ — the variance explained

$$vexp(\mathbf{t}) = \frac{\mathbf{a}^\intercal \mathbf{X}^\intercal \mathbf{X}\mathbf{X}^\intercal \mathbf{X}\mathbf{a}}{\mathbf{a}^\intercal \mathbf{X}^\intercal \mathbf{X}\mathbf{a}}$$

includes part of the variance already explained by the other components. The net increment in total variance explained due to the new component is the variance explained by its residuals orthogonal to the components already in the model. Borrowing from the regression literature, we will refer to the net variance explained as *extra* variance explained, and denote it by *evexp*.

The orthogonal projector onto the columns of a set of components $\mathbf{T}$ is equal to $\mathbf{P_T} = \mathbf{T}(\mathbf{T}^\intercal\mathbf{T})^{-1}\mathbf{T}^\intercal$; the residuals of $\mathbf{X}$ orthogonal to $\mathbf{T}$ are equal to

$$\mathbf{Q} = \mathbf{X} - \mathbf{P_T}\mathbf{X},$$

which we will refer to as *deflated* $\mathbf{X}$ matrix. A new component $\mathbf{t} = \mathbf{X}\mathbf{a}$ deflated of the previous components is equal to $\mathbf{t} - \mathbf{P_T}\mathbf{t} = \mathbf{Q}\mathbf{a}$ and the extra variance that it explains is equal to

$$evexp(\mathbf{t}) = ||\widehat{\mathbf{X}}(\mathbf{Q}\mathbf{a})||^2 = \frac{\mathbf{a}^\intercal \mathbf{Q}^\intercal \mathbf{X}\mathbf{X}^\intercal \mathbf{Q}\mathbf{a}}{\mathbf{a}^\intercal \mathbf{Q}^\intercal \mathbf{Q}\mathbf{a}} = \frac{\mathbf{a}^\intercal \mathbf{X}^\intercal \mathbf{Q}\mathbf{Q}^\intercal \mathbf{X}\mathbf{a}}{\mathbf{a}^\intercal \mathbf{Q}^\intercal \mathbf{Q}\mathbf{a}}. \tag{4}$$

The extra variance explained is equal to the variance explained when $\mathbf{t}$ is the first component (in which case $\mathbf{Q} = \mathbf{X}$) and when it is uncorrelated with all the previously computed components. Since the PCs are mutually uncorrelated, the concept of extra variance explained does not need to be considered in their evaluation.

The extra variance explained is essential for the computation of the sparse components because it measures their contribution accounting for the components already computed. Furthermore, the total variance explained by a set of correlated components $(\mathbf{t}_1, \ldots, \mathbf{t}_c)$ can be expressed as the sum of the extra variances explained as $vexp(\mathbf{t}_1, \ldots, \mathbf{t}_c) = \sum_{j=1}^{c} evexp(\mathbf{t}_j)$. This formulation of the total variance explained is a generalisation of that commonly used in PCA.

A quantity that we will consider later is the variance of $\mathbf{Q}$ explained by

a component. For a component $\mathbf{t} = \mathbf{Xa}$ this is equal to

$$vexp_Q(\mathbf{t}) = ||\widehat{\mathbf{Q}}(\mathbf{t})||^2 = \frac{\mathbf{t}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\mathbf{t}}{\mathbf{t}^\intercal\mathbf{t}} = \frac{\mathbf{a}^\intercal\mathbf{X}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\mathbf{X}\mathbf{a}}{\mathbf{a}^\intercal\mathbf{X}^\intercal\mathbf{X}\mathbf{a}} = evexp(\mathbf{t})\frac{\mathbf{a}^\intercal\mathbf{Q}^\intercal\mathbf{Q}\mathbf{a}}{\mathbf{a}^\intercal\mathbf{X}^\intercal\mathbf{X}\mathbf{a}}.$$

Since $\mathbf{a}^\intercal\mathbf{X}^\intercal\mathbf{X}\mathbf{a} \geq \mathbf{a}^\intercal\mathbf{Q}^\intercal\mathbf{Q}\mathbf{a}$ it follows that

$$vexp_Q(\mathbf{t}) \leq evexp(\mathbf{t}) \leq vexp(\mathbf{t}). \tag{5}$$

$vexp_Q$ will be useful for deriving the sparse components because it is defined by the components expressed in terms of the observed variables (and not of their orthogonal residuals) and is a lower bound for $evexp$.

2.3 Derivation of the projection sparse components

The idea underpinning projection SPCA is to derive the sparse components at each step from a set of variables that explain a given percentage of the (full cardinality) first principal component of the deflated matrix $\mathbf{Q}$, which is the component that explains the most extra variance in that step.

Let $\mathbf{u} = \mathbf{Qv}$ be the first PC of $\mathbf{Q}$, assuming without loss of generality that $\mathbf{v}^\intercal\mathbf{v} = 1$. Since $\mathbf{u}$ is orthogonal to the previous components $vexp(\mathbf{u}) = evexp(\mathbf{u})$. By Equation (2) the (extra) variance explained by the principal component $\mathbf{u}$ simplifies to $vexp(\mathbf{u}) = \mathbf{u}^\intercal\mathbf{u} = \lambda_{\max}$, where $\lambda_{\max}$ is the largest eigenvalue of $\mathbf{Q}^\intercal\mathbf{Q}$. The variance explained by the component $\mathbf{u}$ is an upper bound for the extra variance explained by any component $\mathbf{t} = \mathbf{Xa}$ added to the model, in symbols: $evexp(\mathbf{t}) \leq \lambda_{\max}$.

Assume that the variables in a block $\dot{\mathbf{X}}$ explain a proportion of the variance of $\mathbf{u}$ not less than $0 < \alpha \leq 1$; that is, that $\hat{\mathbf{u}} = \dot{\mathbf{X}}(\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})^{-1}\dot{\mathbf{X}}^\intercal\mathbf{u}$ is such that

$$\frac{\hat{\mathbf{u}}^\intercal\hat{\mathbf{u}}}{\mathbf{u}^\intercal\mathbf{u}} \geq \alpha \text{ or, equivalently, } \hat{\mathbf{u}}^\intercal\hat{\mathbf{u}} \geq \alpha\lambda_{\max}, \tag{6}$$

where the right-hand-side inequality derives from Equation (2). The projection $\hat{\mathbf{u}}$ is a sparse component defined by $\hat{\mathbf{u}} = \dot{\mathbf{X}}\hat{\mathbf{v}}$, with

$$\hat{\mathbf{v}} = (\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})^{-1}\dot{\mathbf{X}}^\intercal\mathbf{u}.$$

The variance of $\mathbf{Q}$ explained by $\hat{\mathbf{u}}$ is equal to:

$$vexp_Q(\hat{\mathbf{u}}) = \frac{\hat{\mathbf{u}}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\hat{\mathbf{u}}}{\hat{\mathbf{u}}^\intercal\hat{\mathbf{u}}} = \frac{\hat{\mathbf{v}}^\intercal\dot{\mathbf{X}}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\dot{\mathbf{X}}\hat{\mathbf{v}}}{\hat{\mathbf{v}}^\intercal\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}}\hat{\mathbf{v}}}. \tag{7}$$

It is evident that if $\hat{\mathbf{u}}$ explains more than a proportion α of the variance of the first PC of $\mathbf{Q}$ it will explain a larger proportion of the variance of the

whole matrix. This can be seen by writing in a singular value decomposition fashion $\mathbf{Q} = \mathbf{u}\mathbf{v}^\intercal + \mathbf{U}_{/1}\mathbf{V}_{/1}{}^\intercal$ (with $\mathbf{u}_j{}^\intercal\mathbf{u}_j = \lambda_j$), where the subscript '/1' denotes that the first column of the matrix has been removed. Then, since $\mathbf{v}^\intercal\mathbf{V}_{/1} = \mathbf{0}$,

$$evexp(\hat{\mathbf{u}}) \geq vexp_Q(\hat{\mathbf{u}}) = \hat{\mathbf{u}}^\intercal\hat{\mathbf{u}} + \frac{\hat{\mathbf{u}}^\intercal\mathbf{U}_{/1}\mathbf{U}_{/1}{}^\intercal\hat{\mathbf{u}}}{\hat{\mathbf{u}}^\intercal\hat{\mathbf{u}}} \geq \hat{\mathbf{u}}^\intercal\hat{\mathbf{u}} \geq \alpha\lambda_{\max}, \tag{8}$$

because of Equations (5) and (6). This proves that the extra variance explained by a projection sparse component $\hat{\mathbf{u}}$ is not less than $\alpha\lambda_{\max}$.

2.4 Connection with least squares SPCA

Merola (2015) derives the LS SPCA components from a block $\dot{\mathbf{X}}$, $\mathbf{t} = \dot{\mathbf{X}}\dot{\mathbf{a}}$, by maximising $vexp_Q(\mathbf{t})$. If the block explains at least a proportion α of the current PC, as in projection SPCA, from Equations (7) and (8) it follows that the extra variance explained by an LS SPCA component is not less than $\alpha\lambda_{\max}$, because

$$evexp(\mathbf{t}) \geq vexp_Q(\mathbf{t}) = \max_a \frac{\dot{\mathbf{a}}^\intercal\dot{\mathbf{X}}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\dot{\mathbf{X}}\dot{\mathbf{a}}}{\dot{\mathbf{a}}^\intercal\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}}\dot{\mathbf{a}}} \geq \frac{\hat{\mathbf{v}}^\intercal\dot{\mathbf{X}}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\dot{\mathbf{X}}\hat{\mathbf{v}}}{\hat{\mathbf{v}}^\intercal\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}}\hat{\mathbf{v}}} \geq \alpha\lambda_{\max}. \tag{9}$$

The LS SPCA loadings are the generalised eigenvectors solution to

$$\dot{\mathbf{X}}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\dot{\mathbf{X}}\dot{\mathbf{a}} = \dot{\mathbf{X}}^\intercal\dot{\mathbf{X}}\dot{\mathbf{a}}\gamma_{\max}, \tag{10}$$

where $\gamma_{\max}$ is the largest generalised eigenvalue and is equal to the variance of $\mathbf{Q}$ explained by the component.

If we consider the orthogonal projection of $\mathbf{Q}$ onto $\dot{\mathbf{X}}$, $\widehat{\mathbf{Q}} = \dot{\mathbf{X}}(\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})^{-1}\dot{\mathbf{X}}^\intercal\mathbf{Q}$, the loadings of the first principal component of $\widehat{\mathbf{Q}}$, $\mathbf{b}$, say, must satisfy:

$$\widehat{\mathbf{Q}}^\intercal\widehat{\mathbf{Q}}\mathbf{b} = \mathbf{Q}^\intercal\dot{\mathbf{X}}(\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})^{-1}\dot{\mathbf{X}}^\intercal\mathbf{Q}\mathbf{b} = \mathbf{b}\gamma_{\max}.$$

Then the vector $\dot{\mathbf{a}} = (\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})^{-1}\dot{\mathbf{X}}^\intercal\mathbf{Q}\mathbf{b}$ satisfies

$$(\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})^{-1}\dot{\mathbf{X}}^\intercal\mathbf{Q}\mathbf{Q}^\intercal\dot{\mathbf{X}}\dot{\mathbf{a}} = \dot{\mathbf{a}}\gamma_{\max}, \tag{11}$$

which is equivalent to Equation (10). Hence the LS SPCA components are the first PCs of the projection of the deflated $\mathbf{Q}$ matrix onto the blocks $\dot{\mathbf{X}}$, while the projection SPCA components are the projections of the first PCs of $\mathbf{Q}$ onto the blocks $\dot{\mathbf{X}}$.

By Equation (9) the variance explained by LS SPCA solutions has a higher lower bound than that explained by projection SPCA components. However, due to the necessity of using greedy approaches to the identification of the blocks of variables, it is not always guaranteed that the LS SPCA components explain more variance than the corresponding projection SPCA components. We leave the comparison of these two SPCA approaches for future research. We can anticipate that in our studies we observed that the extra variance explained is very similar (to the third or fourth decimal figure) for all components.

2.5 Connection with conventional SPCA methods

In conventional methods the SPCA problem is defined by adding sparsity constraints to Hotelling's formulation of PCA in Equation (3). That is, the loadings are the solution to

$$\begin{aligned} \mathbf{a} &= \arg\max_{\mathbf{a}} \mathbf{a}^\intercal \mathbf{S}\mathbf{a} \\ & \quad card(\mathbf{a}) = c,\ \mathbf{a}^\intercal\mathbf{a} = 1. \end{aligned} \tag{12}$$

In some cases the cardinality constraint is substituted with a lasso type penalty $\sum |a_i|$. It is well known that this penalty induces sparsity. Different constraints are added after the first component is computed to avoid trivial solutions.

For most SPCA methods the adoption of model 12 is justified by a not clearly defined *statistical fidelity* criterion. This probably refers to the assumption that under sparsity constraints the norm of the components is equivalent to the variance that they explain. This is not correct, as shown in Section 2.1.

Zou et al. (2006) reduce the minimization of the least squares criterion to the maximisation of the norm of the sparse components by adding an extra constraint to the PCA problem. The least squares PCA problem for a generic component $\mathbf{t} = \mathbf{X}\mathbf{a}$ can be defined as:

$$(\mathbf{a}, \mathbf{p}) = \arg\min_{\mathbf{a},\mathbf{p}} ||\mathbf{X} - \mathbf{X}\mathbf{a}\mathbf{p}^\intercal||^2 \tag{13}$$

where $\mathbf{p}$ is a vector of coefficients. From the first order conditions it follows that $\mathbf{p}^\intercal = (\mathbf{a}^\intercal\mathbf{X}^\intercal\mathbf{X}\mathbf{a})^{-1}\mathbf{a}^\intercal\mathbf{X}^\intercal\mathbf{X}$. Zou et al. (2006) require that $||\mathbf{p}||^2 = \mathbf{p}^\intercal\mathbf{p} = 1$, which implies that $\mathbf{p} = \mathbf{a}$ (so that $\mathbf{p}^\intercal\mathbf{a} = (\mathbf{a}^\intercal\mathbf{X}^\intercal\mathbf{X}\mathbf{a})^{-1}\mathbf{a}^\intercal\mathbf{X}^\intercal\mathbf{X}\mathbf{a} = 1$) and that $||\mathbf{a}|| = 1$. Since $||\mathbf{p}^\intercal\mathbf{p}||^2 = (\mathbf{a}^\intercal\mathbf{S}\mathbf{S}\mathbf{a})/(\mathbf{a}^\intercal\mathbf{S}\mathbf{a})^{-2}$, by the Cauchy-Schwarz inequality the condition $||\mathbf{p}||^2 = 1$ is satisfied if and only if $\mathbf{a} \propto \mathbf{S}\mathbf{a}$, that is, if and only if $\mathbf{a}$ is an eigenvector of $\mathbf{S}$. This condition is satisfied by the

unconstrained PCs but not by the sparse components because when $\mathbf{a}$ is sparse $\mathbf{Sa} \neq \mathbf{a}\lambda$ and $\mathbf{p} = \mathbf{Sa}(\mathbf{a}^\intercal\mathbf{Sa})^{-1} \neq \mathbf{a}$.

The consequence of adding the constraints $\mathbf{p} = \mathbf{a}$ and $\mathbf{a}^\intercal\mathbf{a} = 1$ to the least squares criterion in Equation (13) is that it becomes the maximisation of the norm of the sparse component $\dot{\mathbf{X}}\dot{\mathbf{a}}$ in Equation (12). To see this, assume without loss of generality that the first c variables form the block $\dot{\mathbf{X}}$ and that the block $\overset{\vee}{\mathbf{X}}$ contains the remaining variables, so that $\mathbf{X} = (\dot{\mathbf{X}}, \overset{\vee}{\mathbf{X}})$. Analogously we write $\mathbf{a} = (\dot{\mathbf{a}}^\intercal, \mathbf{0}^\intercal)^\intercal$, where $\mathbf{0}$ is a vector of zeroes of dimension $p - c$. Then it follows that

$$\begin{aligned}
&\arg\min_{\mathbf{a}} ||\mathbf{X} - \mathbf{Xaa}^\intercal||^2 = \arg\min_{\dot{\mathbf{a}}} ||(\dot{\mathbf{X}}, \overset{\vee}{\mathbf{X}}) - \dot{\mathbf{X}}\dot{\mathbf{a}}(\dot{\mathbf{a}}^\intercal, \mathbf{0}^\intercal)||^2 = \\
&\arg\min_{\dot{\mathbf{a}}} ||\dot{\mathbf{X}} - \dot{\mathbf{X}}\dot{\mathbf{a}}\dot{\mathbf{a}}^\intercal||^2 = \arg\max_{\dot{\mathbf{a}}} \dot{\mathbf{a}}^\intercal\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}}\dot{\mathbf{a}} \qquad (14)\\
&\textit{subject to } \dot{\mathbf{a}}^\intercal\dot{\mathbf{a}} = 1.
\end{aligned}$$

The same argument is applicable to Shen and Huang (2008)'s regularised SVD derivation of SPCA.

The maximisation of the norm of the sparse components does not maximise the variance explained by the components and also leads to results that are incompatible with the aim of finding the best sparse approximations of the PCs. This can be can be illustrated by considering the maximisation of the norm of the components on rank deficient data. Such data arise when some of the variables are collinear and when there are fewer observations than variables, which is often the case in gene expression and near infrared data sets.

When $rank(\mathbf{X}) = r < min\{n, p\}$ $p - r$ columns of $\mathbf{X}$ are linearly dependent and any linear combination of the p columns can be expressed in terms of a linearly independent set of r of its columns. Assume without loss of generality that the first r columns of $\mathbf{X}$ are linearly independent and denote them as $\dot{\mathbf{X}}$. Also let $\overset{\vee}{\mathbf{X}}$ be the remaining columns. Then it is simple to prove that

$$\mathbf{X} = [\dot{\mathbf{X}}, \overset{\vee}{\mathbf{X}}] = \dot{\mathbf{X}}\left[\mathbf{I}_r, \dot{\mathbf{X}}^\intercal(\dot{\mathbf{X}}\dot{\mathbf{X}}^\intercal)^-\overset{\vee}{\mathbf{X}}\right] = \dot{\mathbf{X}}\left[\mathbf{I}_r, (\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})^{-1}\dot{\mathbf{X}}^\intercal\overset{\vee}{\mathbf{X}}\right] = \mathbf{P}_{\dot{\mathbf{X}}}\mathbf{X} = \dot{\mathbf{X}}\mathbf{B}, \qquad (15)$$

where the superscript '$-$' denotes the Moore-Penrose generalised inverse and $\mathbf{B} = \left[\mathbf{I}_r, (\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})^{-1}\dot{\mathbf{X}}^\intercal\overset{\vee}{\mathbf{X}}\right]$. Therefore the PCs can be equivalently written as $\mathbf{Xv} \equiv \dot{\mathbf{X}}(\mathbf{Bv}) = \dot{\mathbf{X}}\dot{\mathbf{v}}$, with $\dot{\mathbf{v}} = \mathbf{Bv}$. Since the largest singular value of $\mathbf{X}$ is

not smaller than the largest singular value of $\dot{\mathbf{X}}$, $\dot{\mathbf{v}}^\intercal\dot{\mathbf{v}} \geq 1$. Therefore, the PCs can be equivalently written as sparse components of cardinality r defined by loadings with norm larger than one. However, conventional SPCA methods are prevented from identifying the sparse expression of the PCs because they consider only components with unit norm loadings and the only components with unit norm loadings that are equal to the PCs are the full cardinality PCs themselves. As a consequence, these methods would compute components with a number of variables larger than r. This fact is well documented by several examples available in the SPCA literature. It should be stressed that adding more than r variables to the solution means including redundant variables that are perfectly correlated with the ones already selected, hence unnecessarily increasing its cardinality.

In LS SPCA a component of cardinality equal to the rank of the matrix will be recognised as optimal because by Equation (11) the LS SPCA components are the PCs of $\mathbf{P}_{\dot{\mathbf{X}}}\mathbf{X}$ which, when $rank(\dot{\mathbf{X}}) = r$, is equal to $\mathbf{X}$ by Equation (15). As a matter of fact, LS SPCA components of cardinality larger than the rank of the data matrix cannot be computed because in that case the matrix $\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}}$ would be singular.

As an example of inconsistent results given by conventional SPCA methods we consider a matrix with 100 observations on five perfectly collinear variables defined as

$$x_{ij} = (-1)^i\sqrt{j},\ i = 1, \ldots, 100;\ j = 1, \ldots, 5. \tag{16}$$

The covariance matrix for these variables, $\mathbf{S} = \mathbf{X}^\intercal\mathbf{X}$, has rank one, its largest eigenvalues is equal to 1500 and the others are all zero. The first PC explains all the variance and can be written in terms of any of the variables as $\mathbf{x}_j\sqrt{1500/s_{jj}}$, $j = 1, \ldots, 5$, that is as a cardinality one component with loading larger than one.

The complete set of conventional SPCA solutions of different cardinalities are shown in Table 1. These indicate that x_5 "explains" just 33% of the norm of the PC, x_4 and x_5 together only 60% (with a net increase equal to 27%) and so on. These results mislead the interpretation because they imply that some variables (the ones with larger variance) explain more variance than other collinear ones and that only the full cardinality component (the PC) explains the maximum variance. In reality, any variable explains the same variance as the PC but this cannot be revealed by conventional SPCA methods because the norm of a component with unit norm loadings is bounded by $trace(\dot{\mathbf{X}}^\intercal\dot{\mathbf{X}})$.

Table 1: Optimal conventional SPCA solutions.

	cardinality				
variable	1	2	3	4	5
x_1	0	0	0	0	0.58
x_2	0	0	0	0.60	0.52
x_3	0	0	0.65	0.53	0.45
x_4	0	0.75	0.58	0.46	0.37
x_5	1	0.67	0.50	0.38	0.26
norm	500	900	1200	1400	1500
rel norm	0.33	0.60	0.80	0.93	1.0

The results of applying conventional SPCA to the correlation matrix of the variables in the above example follow a similar pattern. The correlation matrix is a 5×5 matrix of ones with only one non-zero eigenvalue equal to 5. The results of conventional SPCA would be as shown in Table 2 irrespective of which variables are included because the variables are identical. These results lead to the meaningless conclusion that linear combinations of two or more identical variables explain more variance than a linear combination of fewer of them but less than the PC, which is collinear to them.

Table 2: Optimal conventional SPCA solutions for the correlation matrix.

cardinality	1	2	3	4	5
norm	1	2	3	4	5
rel norm	0.2	0.4	0.6	0.8	1.0

Applying LS SPCA to this dataset would yield a single cardinality one component both for the covariance and the correlation matrices, as required. This can be seen by considering that the variance explained by any variable x_j is equal to

$$vexp(\mathbf{x}_j) = \frac{\mathbf{x}_j{}^\intercal \mathbf{X}\mathbf{X}^\intercal \mathbf{x}_j}{\mathbf{x}_j{}^\intercal \mathbf{x}_j} = \frac{\sum_{i=1}^{5} s_{ij}^2}{s_{jj}} = \sum_{i=1}^{5} s_{ii} = tr(S),$$

because $cor(x_i, x_j)^2 = 1 = s_{ij}^2/(s_{ii}s_{jj})$. Which variable is chosen depends on the algorithm used. An analogous example for blocks of collinear variables can be derived using the artificial data example introduced in Zou et al. (2006), as illustrated in Merola (2015).

The case of collinear variables above is useful also to illustrate the fundamental difference between conventional SPCA and LS SPCA. The maximisation of the norm identifies blocks of highly correlated variables, which give similar information; the maximisation of the variance explained selects less

correlated variables which are important for explaining the variance also of the other variables in the set. For this reason, LS SPCA yield lower cardinality components that are better approximations of the PCS. Merola (2015) also discusses other unconvincing aspects of conventional SPCA methods.

3 Computational details

The basic projection SPCA algorithm is outlined in Algorithm 1. The algorithm is straightforward and simple to implement. Its complexity depends on how the variables for each block are selected (line 10) and the eigenvectors are computed (line 9).

Determining the variables that explain a given percentage of the current PC is a problem analogous to model selection in multiple regression and feature selection in machine learning. These problems can be solved with different greedy approaches and algorithms. The question of which algorithm performs best is a long-standing problem which is difficult to answer because of the existence of competing desired objectives. A recent literature review can be found in Kumar and Minz (2014). Many of these algorithms can be used for selecting the variables for computing projection SPCA.

In our implementation we chose the simple forward selection used in multiple regression analysis for maximising the coefficient of determination because it is computationally very efficient. In this algorithm the variables that sequentially explain the most extra variance of the current PC enter the solution until a predetermined threshold for the coefficient of determination is reached. This method can be seen as a QR decomposition with supervised pivoting and can be implemented efficiently using updating formulae (e.g. see Björck, 2015).

Since for each component only the first eigenvector of the deflated covariance matrix $\mathbf{Q}_j{}^\intercal\mathbf{Q}_j$ is required, these can be computed economically using the power method. The power method computes approximated results iteratively and is has complexity about $O(p^2)$, while algorithms that compute the whole set of eigenvectors accurately are about $O(p^3)$. The power method is used in extremely high dimensional problems such as, for example, Google page ranking (Bryan and Leise, 2006). The algorithm can be modified to compute LS SPCA components by changing line 11 accordingly and modifying the power method for the computation of generalised eigenvectors.

The theoretical time complexity of the algorithm cannot be computed exactly but we expect the complexity of the whole algorithm to be less than $O(p^3)$ because it does not contain operations of this complexity or higher. In the next section we will analyse the time taken on the datasets empirically.

Algorithm 1 Projection SPCA algorithm

procedure PSPCA($\mathbf{X}$, α, *stopRuleCompute*)
  initialize
    $\mathbf{Q} \leftarrow \mathbf{X}$
    $j \leftarrow 0$
    stopCompute $\leftarrow$ FALSE
  end initialize
  while stopCompute = FALSE **do** ▷ **start components computation**
    $j \leftarrow j + 1$
    $\mathbf{u}_j : \mathbf{Q}\mathbf{Q}^\intercal \mathbf{u}_j = \mathbf{u}_j \lambda_j$ ▷ compute first PC of $\mathbf{Q}$
    $ind_j \leftarrow \{i_1, \ldots, i_{c_j}\} : ||\hat{\mathbf{u}}_j||^2 \geq \alpha\lambda_j$ ▷ variables selection output
    $\mathbf{a}_j \leftarrow (\dot{\mathbf{X}}_j{}^\intercal \dot{\mathbf{X}}_j)^{-1} \dot{\mathbf{X}}_j{}^\intercal \mathbf{u}_j$ ▷ j-th sparse loadings, $\dot{\mathbf{X}}_j$ are columns of $\mathbf{X}$ in ind_j
    $\mathbf{t}_j \leftarrow \dot{\mathbf{X}}_j \mathbf{a}_j$ ▷ j-th sparse component
    if stopRuleCompute = FALSE **then** ▷ stop rule on total variance explained or number of components
      $\mathbf{Q} \leftarrow \mathbf{Q} - \frac{\mathbf{t}_j \mathbf{t}_j{}^\intercal \mathbf{Q}}{\mathbf{t}_j{}^\intercal \mathbf{t}_j}$ ▷ deflate $\mathbf{X}$ of current component
    else
      stopCompute $\leftarrow$ TRUE ▷ **terminate components computation**
    end if
  end while
end procedure

4 Numerical examples

In this section we will give an example of data analysis using projection SPCA and illustrate the computational times and sparse reduction attained on large datasets. For simplicity of exposition, we refer to the size of the loadings meaning the absolute value of the non-zero ones.

4.1 Sparse principal components analysis of crime data

As an illustration of sparse principal component analysis we compare the results of applying simple thresholding and projection SPCA to the correlation matrix of a dataset containing socioeconomic measures collected in the 1990s for different US cities (available at the UCI repository, Frank and Asuncion, 2010).The data were originally used to explain the rate of violent crime in the different cities. We deleted 23 variables with missing values. The final set contains 1994 observations on 99 variables. The example is given only to illustrate an application of sparse principal components analysis and does not claim to be a deep analysis of these data, which is beyond the scope of this paper and our expertise.

The ordered loadings of the first and second PCs scaled to percentage

contributions that is, scaled to unit sum of absolute values are shown in Figure 1. Clearly, it is difficult to choose a threshold value to separate significant and nonsignificant loadings as they decrease at a regular rate.

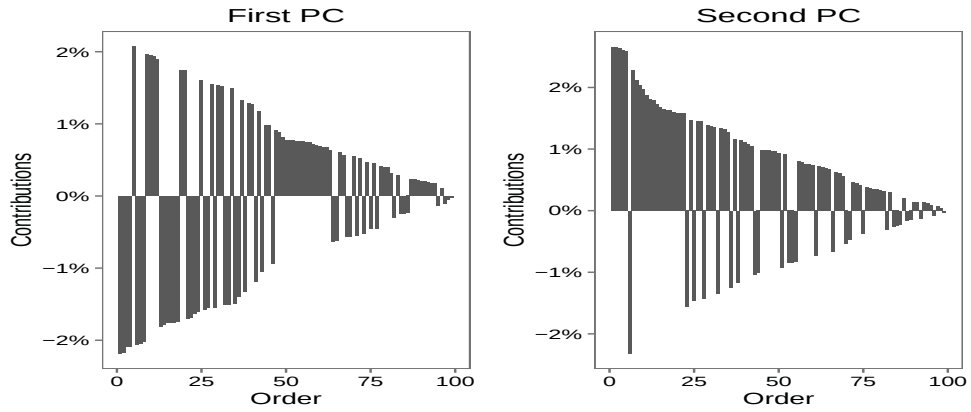

Figure 1: Ordered contributions of the first two PCs of the crime dataset.

The summary results of projection SPCA with $\alpha = 95\%$ are shown in Table 3. In this table the relative cumulative variance explained (rCvexp) denotes the ratio of the cumulative variance explained by the components over that explained by the PCs. This is, it is computed as

$$\text{rCvexp}(\hat{\mathbf{u}}_1, \ldots, \hat{\mathbf{u}}_c) = \frac{vexp(\hat{\mathbf{u}}_1, \ldots, \hat{\mathbf{u}}_c)}{vexp(\mathbf{u}_1, \ldots, \mathbf{u}_c)} = \frac{\sum_{j=1}^{c} evexp(\hat{\mathbf{u}}_j)}{\sum_{j=1}^{c} vexp(\mathbf{u}_j)}.$$

As expected, rCvexp is always higher than 95% and the cardinality of the sparse components is quite low. The first two sparse components are highly correlated with the corresponding PCs while the correlation decreases for successive components. The R^2 statistics refer to the regression of the logarithm of crime rate on the components computed thus far.

Table 3: Summary statistics for the first five projection SPCA components.

	Comp1	Comp2	Comp3	Comp4	Comp5
% Cumulative vexp	24.4	40.8	49.8	57.2	62.7
% Relative cumulative vexp	96.5	96.5	96.5	96.6	96.6
Cardinality	3	5	7	9	8
Correlation with PC	0.97	0.98	0.94	0.83	0.92
% R^2 log(crime rate) on PCs	44.0	49.0	50.9	51.3	53.3
% R^2 log(crime rate) on sparse comp.	45.7	51.9	56.4	57.2	57.2

Table 4 shows the ten largest contributions of the first PC and the three non-zero contributions of the first projection sparse component. The last

column shows the coefficients of determination resulting from regressing each variable on all other variables in the set (squared multiple correlations), R_j^2.

The ten largest PC contributions assign about the same weight to different measures of two indicators: income and the presence of both parents in a household. The corresponding variables are highly correlated among themselves, as shown by the large R_j^2 values. The sparse contributions assign about 50% of weight to income, about 37% to families with two parents, and 12% to a new indicator: large family size. This last variable is hardly correlated with the others and its contribution to the PC ranks 46th in size. Combining just three variables out of 99, the sparse component explains 96.5% of the variance explained by the first PC, indicates the key features of the PC and identifies another important variable which would have been difficult to find by analysing the PC loadings.

Table 4: Largest ten contributions of the first full cardinality PC and those of the first projection SPCA component.

Contribution	Variable	R_j^2
PCA		
-2.2%	**median family income (differs from household income)**	0.98
-2.2%	median household income	0.98
-2.1%	**percentage of children in family housing with two parents**	0.98
-2.1%	percentage of households with investment/rent income in 1989	0.83
2.1%	percentage of people below the poverty level	0.84
-2.1%	percentage of families (with children) with two parents	0.98
-2.0%	percentage of children aged 4 and under in two-parent households	0.90
-2.0%	per capita income	0.92
2.0%	percentage of households with public assistance income in 1989	0.75
2.0%	percentage of occupied housing units without a phone (in 1990, this was rare!)	0.76
Proj SPCA		
-51%	median family income (differs from household income)	0.51
-37%	percentage of children in family housing with two parents	0.52
12%	**percentage of family households that are large (6 or more)**	0.09

The ten largest contributions of the second PC, shown in Table 5, assign about 92% of the weights (rescaled) to highly correlated variables regarding immigrants: the percentage of immigrants by time of arrival and skills in English language, and 8% to living in dense housing conditions. Instead, the five sparse contributions, also shown in Table 5, assign 52% of the weights to one variable regarding immigrants, while the remaining weights are evenly divided into four variables concerning different social aspects such as the percentage of older residents and the percentage of people living in their own house. These variables have low multiple correlation. The indicator for large families, which was also present in the first component, has high correlation with the indicator for dense housing, which is an important contributor to

the PC. The positive contribution of the percentage of people living in their own house is probably due to the cost of paying a mortgage over household income and to the fact that older residents are more likely to own their house.

Table 5: Largest ten contributions of the second full cardinality PC and those of the second sparse components computed with projection SPCA.

Contribution	Variable	R_j^2
PCA		
2.7%	**percentage of population who immigrated within the last 10 years**	1.00
2.7%	percentage of population who immigrated within the last 8 years	1.00
2.6%	percentage of population who immigrated within the last 5 years	1.00
2.6%	percentage of population who immigrated within the last 3 years	0.98
2.6%	percentage of foreign-born people	0.96
-2.3%	percentage of people who speak only English	0.95
2.3%	percentage of people who do not speak English well	0.94
2.1%	**percentage of people in dense housing (more than 1 person per room)**	**0.86**
2.0%	percentage of population of Asian heritage	0.63
2.0%	percentage of population of Hispanic heritage	0.90
Proj SPCA		
42%	percentage of population who immigrated within the last 10 years	0.11
-15%	**percentage of population aged 65 and over**	0.06
15%	**owner-occupied housing - upper quartile value**	0.55
14%	**percentage of family households that are large (6 or more)**	0.47
13%	**number of people living in areas classified as urban**	0.32

In summary, in this example the sparse components combining few key variables explain over 96% of the variance explained by the PCs. In contrast, large PC loadings correspond to clusters of mutually highly correlated variables which shadow other important but less correlated variables.

The sparse components computed with LS SPCA are often better at predicting variables not included in the dataset than the PCs. This probably depends on the sparse components giving higher weight to variables that are less correlated with others in the set, but which may be more correlated with an exogenous response. This is not guaranteed for all applications but, in this example, the sparse components yield higher coefficients of determination for the logarithms of crime rate than the PCs, as shown in Table 3. These coefficients of determination are quite high considering that the coefficient of determination for regressing the logarithm of crime rate on all 99 variables is 69.1%.

4.2 Computational results for different datasets

In this section we report the results of running 20 repetitions of projection SPCA on datasets with different numbers of variables, from less than 100 to over 16,000. For all datasets we required that each sparse component

explained at least 95% of the corresponding PC. The computational times reported were measured using an implementation of the algorithm in `C++` embedded in `R` using the packages Rcpp and RcppEigen.

It is well known that `R` is an inefficient language (Morandat et al., 2012), so this choice makes the program run slower than if it were written only in `C++` but has other advantages. Therefore, the times shown are only indicative of the performance of the algorithms on large matrices, which can be improved if necessary.

The datasets that we consider in this section, listed in Table 6, have been taken from various sources, mostly from the data distributed with the Elements of Statistical Learning book (ESL) (Hastie et al., 2009) website. Other sets were taken from the UCI Machine Learning Repository (Frank and Asuncion, 2010). The remaining sets were taken from different sources, see Table 6 for details. Most of the datasets contain gene expression data. The characteristic of these sets is to have a large number of features and fewer objects. The largest dataset (Ramaswamy et al., 2001) has been used to test other SPCA methods (Zou et al., 2006; Sriperumbudur et al., 2009; Wang and Wu, 2012).

Table 6: Description of the datasets used for numerical comparison.

Name	Samples	Features	Description	Source
Crime	1994	99	social data	UCI Repository[a]
Isolet	6238	716	character recognition	UCI Repository[b]
Ross NCI60	60	1375	gene expression	`R` package made4 [c]
Khanh	88	2318	gene expression	ESL
Phoneme	256	4509	speech recognition	ESL
NCI60	60	6830	gene expression	ESL
Protein	11	7466	protein cytometry	ESL
Radiation	58	12625	gene expression	ESL
Ramaswamy	144	16063	gene expression	Broadinstitute repository[d]

[a] `https://archive.ics.uci.edu/ml/datasets/Communities+and+Crime`
[b] `https://archive.ics.uci.edu/ml/datasets/ISOLET`
[c] `http://bioconductor.org/packages/release/bioc/html/made4.html`
[d] `http://software.broadinstitute.org/cancer/software/genepattern/datasets`

Tables 7 and 8 show the median computational time, over 20 repetitions, taken to compute ten sparse components for nine datasets considered, together with the cardinalities. Table 7 also shows the relative cumulative variance explained, which is always larger than 95% as required. The same is true for larger datasets (not shown). The cardinalities are also very low for the largest datasets. Such dramatic reduction may be partly due to the fact that these sets have fewer observations than variables. The times taken to compute the first component vary from four milliseconds to 2.7 seconds. Computing ten components for the largest set takes less than one and a half

minutes.

Computing the theoretical time complexity of the projection SPCA algorithm is not possible because the convergence times for both the forward selection and power method algorithms depend on the data structure. Since the power method is at least $O(n^2)$ and there are no operations $O(n^3)$, we expect the computational complexity to be within these two limits. The time taken to compute additional components grows almost linearly for fixed dataset size, and the time taken to compute a fixed number of components increases non-linearly with the size of the datasets.

Table 7: Median computational times, relative cumulative variance explained, and cardinalities of the first 10 projection SPCA components with $\alpha = 0.95$ for three medium-size datasets. Time is expressed in milliseconds.

Name	Isolets			Ross NCI60			Khanh		
No. vars	716			1375			2318		
No. Components	Time	rCvexp	card	Time	rCvexp	card	Time	rCvexp	card
1 Comp	4.0	95.4%	5	24.1	96.8%	5	126.2	96.4%	6
2 Comps	15.2	95.4%	17	65.1	96.6%	6	260.6	96.2%	6
3 Comps	26.9	95.5%	15	102.9	96.7%	6	360.8	96.4%	4
4 Comps	45.6	95.5%	26	155.0	96.7%	9	555.6	96.4%	8
5 Comps	63.34	95.5%	27	219.6	96.7%	9	731.3	96.5%	9
6 Comps	80.71	95.6%	25	281.4	96.8%	9	969.8	96.5%	8
7 Comps	103.0	95.6%	32	343.1	96.8%	11	1167.0	96.5%	10
8 Comps	125.5	95.8%	36	418.1	96.9%	11	1447.0	96.6%	12
9 Comps	148.7	95.9%	32	499.5	96.9%	12	1640.0	96.6%	10
10 Comps	175.2	95.9%	42	580.0	97.0%	13	1841.0	96.7%	12

Table 8: Median computational times and cardinalities of the first 10 projection SPCA components with $\alpha = 0.95$ for five large-size datasets. Time is expressed in seconds.

Name	Phoneme		NCI60		Protein		Radiation		Ramaswamy	
No. vars	4509		6830		7466		12625		16063	
No. comp.	Time	card	Time	card	Time	card	Time	card	Time	card
1 Comp	0.1	1	0.6	4	0.4	1	2.7	5	2.3	3
2 Comps	0.4	4	1.9	5	1.1	2	6.1	5	6.9	4
3 Comps	1.1	13	2.9	6	1.8	1	12.8	7	12.5	6
4 Comps	2.0	16	4.5	8	2.4	2	17.7	7	19.5	8
5 Comps	3.3	20	6.7	10	3.3	2	23.7	7	28.0	9
6 Comps	4.4	23	8.2	9	4.2	3	30.3	8	37.6	8
7 Comps	5.8	27	10.0	8	5.6	3	37.3	8	46.2	10
8 Comps	7.0	28	12.0	10	6.5	3	44.6	9	60.4	13
9 Comps	8.7	27	14.2	10	7.4	4	52.0	9	73.0	17
10 Comps	10.6	26	16.4	10	8.2	4	59.0	8	85.6	16

To estimate the time complexity function for c components and p variables, $T(c, p)$, we assume that it is polynomial in both parameters, which can be written as

$$T(c,p) = kd^{\alpha}p^{\beta}\epsilon \text{ or } \log(T(c,p)) = \log(k) + \alpha\log(c) + \beta\log(p) + \eta.$$

The fit obtained with the data available is almost perfect with $R^2 = 0.99$, as shown in Table 9. As expected, the estimated coefficients show that computational time increases slightly faster than linearly with the number of components computed while it is approximately $O(p^2)$ with respect to the dataset size. It should be noted that the time taken for matrices with fewer objects than features can be reduced by using the inverse svd. We do not report these times because are peculiar to the matrices while the ones reported refer generically to matrices with a given number of variables.

Table 9: Output of the regression of $\log(T(c,p))$ on $\log(c)$ and $\log(p)$.

Coefficients	Estimate	Std. error	t-value	Pr(> |t|)	
$\log_{10}(k)$	-5.06	0.08	-61.44	< 2e-16	***
α	1.46	0.05	30.81	< 2e-16	***
β	2.03	0.02	93.37	< 2e-16	***

Residual standard error: 0.1362 on 87 degrees of freedom
Multiple R-squared: 0.9911, Adjusted R-squared: 0.9909
F-statistic: 4833 on 2 and 87 DF, p-value: < 2.2e-16

4.3 Comparison with conventional SPCA

In this section we compare the performance of the first sparse components computed with a conventional SPCA method and with projection SPCA. As conventional SPCA method we used SPCA-IE (Wang and Wu, 2012) with *amvl* criterion. This method was shown to perform similarly to other SPCA methods, it does not require to choose arbitrary sparsity parameters and is simple to implement.

We compare the norm of the components, the relative cumulative variance that they explain and their correlation with the full PC for different cardinalities. Since the results for large matrices are quite similar we limit this comparison to four datasets: Crime, Isolets, Khanh and Ramaswamy. For the last dataset we computed the conventional sparse components using simple thresholding, stopping the computation to cardinality 200; details of the performance of components with larger cardinality computed with different conventional SPCA methods can be found in the papers cited above.

Figure 2 compares the performance of the components computed with different methods. The LS SPCA values for the rank deficient Khanh and Ramaswamy datasets are available only until the solutions reach full rank cardinality (87 and 143, respectively) at which the components explain the maximum possible variance. Clearly SPCA-IE outperforms projection SPCA in the norm of the components. However, the latter method guarantees higher variance explained and closer convergence to the PC with much lower

cardinality.

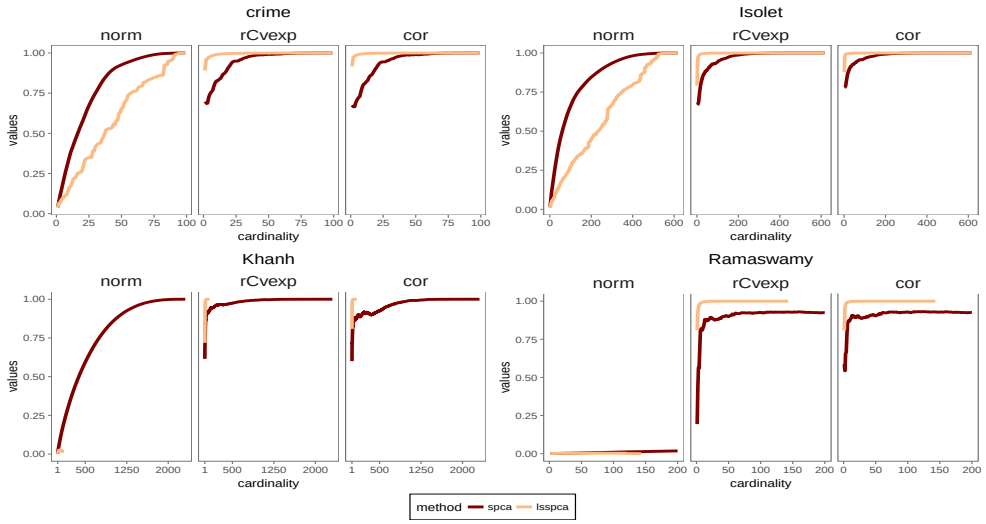

Figure 2: Norm, rCvexp and correlation with the PCs versus cardinality for the first sparse components computed with SPCA-IE (spca) and projection SPCA (lsspca).

This phenomenon is more evident for large rank deficient datasets when conventional sparse components combination of hundreds of variables explain less variance than projection components which are combinations of much fewer variables. As explained in Section 4.3, projection SPCA components reach the maximum when their cardinality is equal to the matrix rank while conventional SPCA components reach the maximum when their cardinality is equal to the number of variables. The plots show clearly that the components' norms are not related to the variance that they explain or to their correlation with the PC, confirming the theoretical conclusions given above.

Table 2 shows the cardinality with which the components computed with the two methods reached 99.9% rCvexp (or 0.998 correlation with the PC). For the Ramaswamy dataset the threshold for the variance explained is not reached with cardinality 200. The last column of the table shows rCvexp for the conventional SPCA components with the same cardinality at which the projection SPCA components reach 99.9% rCvexp. Again, it can be appreciated the much greater efficiency of projection SPCA at approximating the PCs.

Table 10: Cardinality needed to reach 99.9% rCvexp by the components computed with projection SPCA, second column, and SPCA-IE, third column. The fourth column shows the rCvexp for the SPCA-IE components with the same cardinality as the corresponding projection SPCA component.

Dataset	Rank	Proj. SPCA	SPCA	rCvexp SPCA
Crime	99	38	74	85.200%
Isolet	716	123	439	59.600%
Khanh	87	28	1338	0.060%
Ramaswamy	143	23	> 200	0.002%

5 Discussion

Conventional SPCA method have been shown to give results similar to simple thresholding if not worse (Zou et al., 2006; Wang and Wu, 2012). Thresholding has been proven to give misleading results (e.g. Cadima and Jolliffe, 1995). Since the loadings are proportional to the covariances of the variables with the PCs, the largest loadings correspond to variables that are highly correlated with the current PC and among themselves. These sets of variables are not very informative because they contain different measures of the same features.

Zou et al. (2006) propose three properties of a good SPCA method:

- Without any sparsity constraint, the method should reduce to PCA.
- It should be computationally efficient for both small p and big p data.
- It should avoid misidentifying the important variables.

In the light of what explained above, the first property is not enough for a good method. The second is not necessary for the most commonly analysed datasets and the third one is vague because importance is not defined and variables known to be unimportant could be directly eliminated from the analysis.

We suggest the following properties for a good SPCA method:

- Without any sparsity constraint the method should reduce to PCA.
- It should identify the sparsest expression of the principal components.
- The addition of a variable perfectly correlated with one or more variables already in the solution should not improve the objective function.

The last property eliminates redundant variables from the solution and should deter the inclusion of highly correlated ones. Conventional SPCA methods

do not have the last two properties while methods based on LS SPCA do. It is surely possible that different methods would have these properties.

Projection SPCA methodology is intuitive and can be understood by users who do not have a deep knowledge of numerical optimisation. The only parameter to be set for computing the solutions is the proportion of variance explained, the meaning of which is also easily understandable. The algorithm is simple to implement and scalable to large datasets. Most conventional SPCA methods are based on complicated numerical optimisations and require setting values of parameters the effect of which is difficult to understand.

Different variable selection methods could be used for projection SPCA. Future research could explore the effect of using different selection methods, such as least angle or lasso regression, for example.

In this work we have developed a framework for computing LS SPCA solutions. These solutions give sparse components that closely approximate the full PCs with low cardinality. These can be efficiently computed also for very large datasets. We also show that conventional SPCA methods suffer from a number a drawbacks which make the solution provided less attractive than the corresponding LS SPCA solutions.

LS SPCA is implemented in the R package *spca* available on GitHub. Projection SPCA will be added to this package in the future.

References

References

Björck, Å., 2015. Numerical Methods in Matrix Computations. Vol. 59 of Texts in Applied Mathematics. Springer International Publishing.

Bryan, K., Leise, T., 2006. The $25,000,000,000 eigenvector: The linear algebra behind Google. SIAM REVIEW 48 (3), 569–581.

Cadima, J., Jolliffe, I. T., 1995. Loadings and correlations in the interpretation of principal components. Journal of Applied Statistics 22 (2), 203–214.

Frank, A., Asuncion, A., 2010. UCI machine learning repository. University of California, Irvine, School of Information and Computer Sciences, isolet dataset available at https://archive.ics.uci.edu/ml/datasets/ISOLET.

Hastie, T., Tibshirani, R., Friedman, J., 2009. The elements of statistical learning: data mining, inference and prediction, 2nd Edition. Springer.

Hotelling, H., 1933. Analysis of a Complex of Statistical Variables with Principal Components. Journal of Educational Psychology 24, 498–520.

Jeffers, J., 1967. Two case studies in the application of principal component. Applied Statistics 16, 225–236.

Jolliffe, I., Uddin, M., 2000. The simplified component technique: An alternative to rotated principal components. Journal of Computational and Graphical Statistics 9 (4), 689–710.

Kumar, V., Minz, S., 2014. Feature selection: A literature review. Smart CR 4 (3), 211–229.

Merola, G., 2015. Least squares sparse principal component analysis: a backward elimination approach to attain large loadings. Australia & New Zealand Journal of Statistics 57, 391–429.

Moghaddam, B., Weiss, Y., Avidan, S., 2006. Spectral bounds for sparse pca: Exact and greedy algorithms. In: Advances in Neural Information Processing Systems. MIT Press, pp. 915–922.

Morandat, F., Hill, B., Osvald, L., Vitek, J., 2012. Evaluating the design of the r language: Objects and functions for data analysis. In: Proceedings of the 26th European Conference on Object-Oriented Programming. ECOOP'12. Springer-Verlag, Berlin, Heidelberg, pp. 104–131.

Pearson, K., 1901. On lines and planes of closest fit to systems of points in space. Philosophical Magazine 2 (6), 559–572.

Ramaswamy, S., Tamayo, P., Rifkin, R., Mukherjee, S., Yeang, C. H., Angelo, M., Ladd, C., Reich, M., Latulippe, E., Mesirov, J. P., Poggio, T., Gerald, W., Loda, M., Lander, E. S., Golub, T. R., dec 2001. Multiclass cancer diagnosis using tumor gene expression signatures. Proceedings of the National Academy of Sciences of the United States of America 98 (26), 15149–15154.

Shen, H., Huang, J., 2008. Sparse principal component analysis via regularized low rank matrix approximation. Journal of Multivariate Analysis 101, 10151034.

Sriperumbudur, B., Torres, D., Lanckriet, G., 2009. A D.C. programming approach to the sparse generalized eigenvalue problem. Computer Engineering 1 (1), 40.

Wang, Y., Wu, Q., 2012. Sparse pca by iterative elimination algorithm. Advances in Computational Mathematics 36, 137–151.

Zou, H., Hastie, T., Tibshirani, R., 2006. Sparse principal component analysis. Journal of Computational and Graphical Statistics 15 (2), 265–286.